# Write Where It Matters: Policy-Guided Watermarks for 3D Gaussian Splatting

## Supplementary Material

## 7. Write Where It Matters

We further explicitly visualize $\mathbf{W}^{(k)}$ by replacing each anchor's opacity with its $[0, 1]$ action, setting constant red, and rendering the scene. As shown in Fig. 5, high weights cluster in structured, high-frequency, geometrically meaningful regions, while large smooth surfaces receive low weights. Compared to high uncertainty regions (anchors above average uncertainty), W2M does not simply "write on uncertainty", but targets information-rich structures while avoiding overly unstable regions.

Additionally, we compare full W2M against two variants that mask the policy weights (Top-only: only anchors with $w \geq 0.8$ are updated; Top-disabled: those with $w \geq 0.8$ are frozen) on LLFF with 48 bits. As shown in Table 6, disabling high-$w$ anchors increases PSNR but ***sharply*** reduces bit accuracy, while writing only on high-$w$ anchors preserves most of the accuracy with minor fidelity change. This directly tests whether the regions assigned high weights are indeed the ones that carry watermark information.

Table 6. Ablation: high-weight anchors matter.

| | W2M (Full) | Top $\geq 0.8$ Only | Top $\geq 0.8$ Disabled |
|---|---|---|---|
| PSNR↑ | 38.10 | 39.17 | 42.02 |
| Acc.↑ | 98.75 | 95.83 | 62.77 |

## 8. Qualitative Results

Figures 6–8 provide per-scene visual examples for LLFF, Mip-NeRF 360, and Blender in 48 bit. Each row shows the original render, our watermarked render, the raw difference, and an amplified difference ($\times 10$) for better visualization.

Across datasets, the residuals remain low-magnitude and spatially sparse, matching the anchor-wise policy design that "writes where it matters" for decoding while preserving visual fidelity. In LLFF (Fig. 6), Flower and Leaves maintain leaf venation and fine textures; the amplified maps mainly trace occlusion boundaries and edgelets. Room shows light-structure transitions preserved with differences confined to contour lines. In Mip-NeRF 360 (Fig. 7), Bicycle concentrates signal along spokes and rim edges; Kitchen places edits on tile/texture seams; Stump localizes around foliage boundaries. In Blender (Fig. 8), Drums and Lego highlight edges and specular rims, while Mic uses filament-like details; Ship exhibits sparse yet structured edits along rigging and masts.

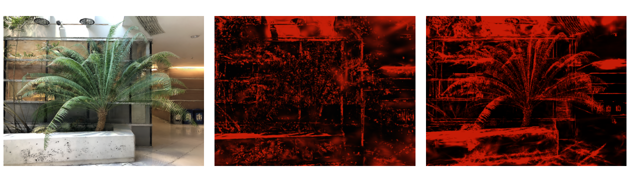

Figure 5. Clean render (left), W2M heatmap (middle), high uncertainty region (right).

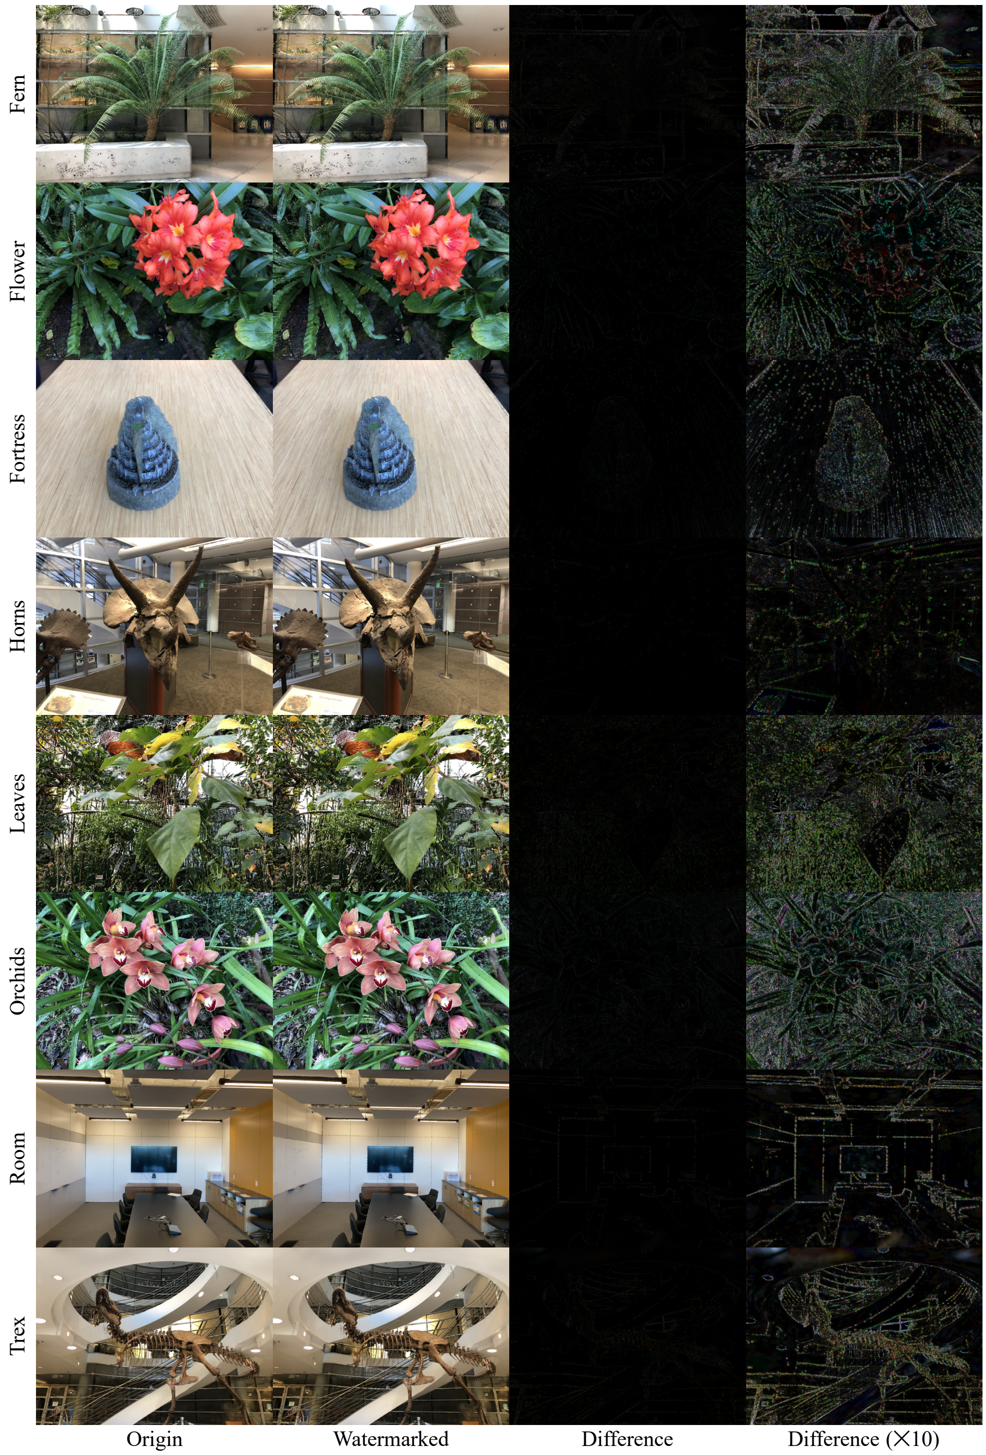


Figure 6. Qualitative results of the proposed W2M on the LLFF datasets. In each line, we present the original rendered image, the watermarked rendered image, the difference map between the watermarked and original rendered images, and an amplified difference map (×10) for better visualization.


Figure 7. Qualitative results of the proposed W2M on the Mip-NeRF 360 datasets. In each line, we present the original rendered image, the watermarked rendered image, the difference map between the watermarked and original rendered images, and an amplified difference map (×10) for better visualization.

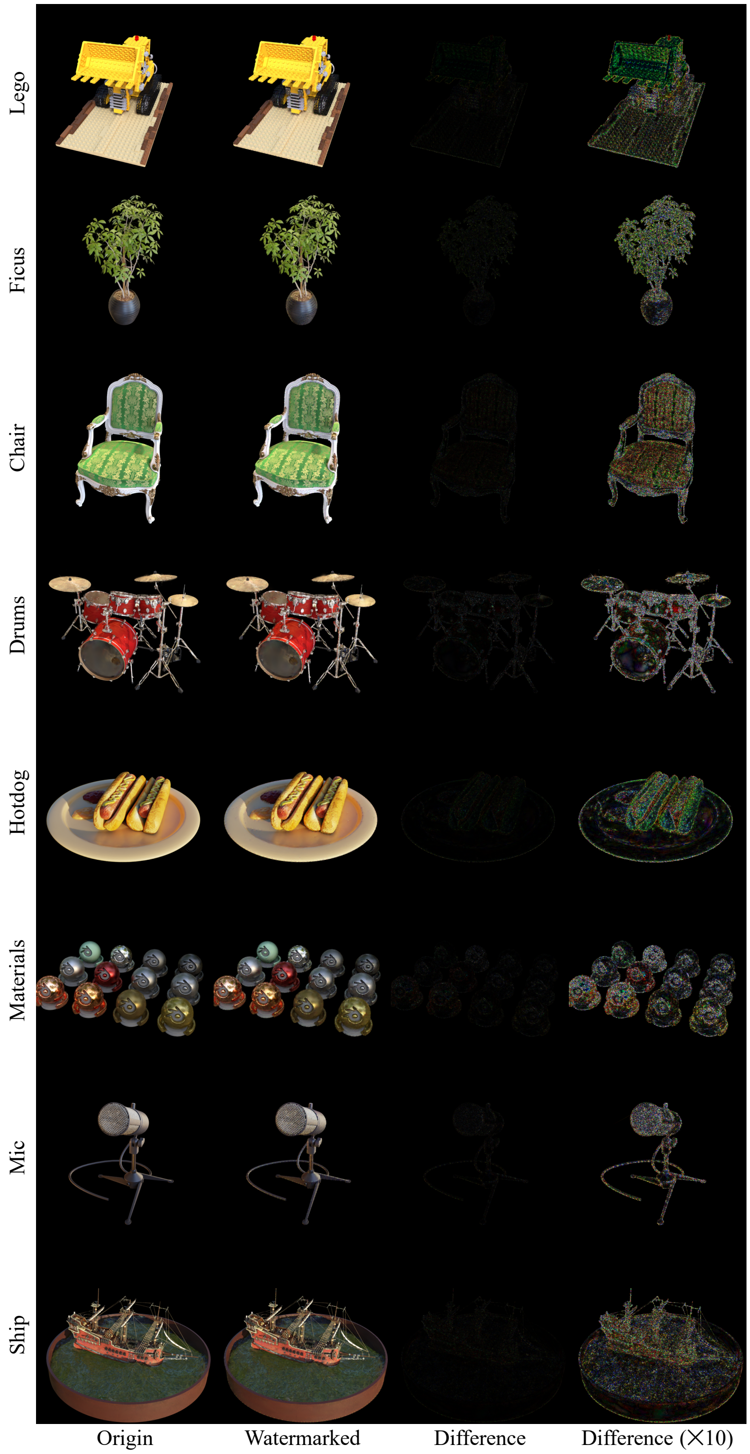


Figure 8. Qualitative results of the proposed W2M on the Blender datasets. In each line, we present the original rendered image, the watermarked rendered image, the difference map between the watermarked and original rendered images, and an amplified difference map (×10) for better visualization.